



UNITED STATES PATENT AND TRADEMARK OFFICE

ser

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/812,200

03/29/2004

James A. Mott

SUN-P8925

2866

57913 7590 08/09/2007
SUN MICROSYSTEMS, INC.
c/o PARK VAUGHAN & FLEMING, LLP
2820 FIFTH STREET
DAVIS, CA 95618

EXAMINER

VU, THONG H

ART UNIT

PAPER NUMBER

2616

MAIL DATE

DELIVERY MODE

08/09/2007

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/812,200.

Applicant(s)

MOTT, JAMES A.

Examiner

Thong H. Vu

Art Unit

2616

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 29 March 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-45 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-45 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

1. Claims 1-45 are pending.
2. The co-pending application 10/811642 file 3/29/2004.

Claim Rejections - 35 USC § 103

Claims 1-8,12-20,23-30,34-38,42-45 are rejected under 35 U.S.C. 103(a) as being unpatentable over Brian et al [Brian 2005/0129039 A1] in view of Payson [7,194,661 B1].

3. As per claim 1, Brian discloses In a communication device coupled to an InfiniBand network and an external communication system, a method of transferring a communication from the InfiniBand network to the external communication system, the method comprising:

receiving packets for each of multiple queue pairs terminating at the communication device [Brian, multiple queue pairs, Fig 3; termination, 0024];

for each of said queue pairs [Brian, queue pair, 0019]:

reassembling in a shared memory contents of said packets into communications to be transmitted to the external communication system [Brian, reassemble buffers 34, Fig 3; shared queue, 0142,0147]; and

Brian also taught a list of send queues [Brian, Fig 1A, 0023] and a list of physical pages associated with memory region [Brian, 0025]. However Brian does not explicitly detail

maintaining an associated linked list identifying locations in said shared memory in which said communications are reassembled; and

when a communication is reassembled for a first queue pair, identifying to a transmission module a portion of a first linked list associated with said first queue pair, wherein said first linked list portion identifies shared memory locations in which said communication was reassembled.

In the same endeavor, Payson discloses an Infiniband architecture system area network including a packet context manager manages reassembly of cells into a packet context using storage memory and free list/link list memory for transmission, col 7 lines 48-60; wherein the memory is shared between ingress and egress data path, col 12 lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with shared memory and identifying the queue pair to the link list for transmission as taught by Payson into the Brian's apparatus in order to utilize the RDMA interface controller. Doing so would provide an assist that detects when there may be an actual error in the network switch so that the fault can be corrected to thereby reduce the possibility of multiple faults occurring at the same time.

4. As per claim 2, Brian-Payson disclose only after said reassembled communication is transmitted, removing said portion of said first linked list from said first linked list [Payson, link list, col 7 lines 48-60].

5. As per claim 3, Brian-Payson disclose said reassembling comprises: as said packets are received from the InfiniBand network, queuing said contents directly into

said shared memory; wherein said shared memory serves as receive queues for each of said multiple queue pairs [Brian, placed directly by OutLogic, 0065].

6. As per claim 4, Brian-Payson disclose transmitting said reassembled communication from said shared memory; wherein said shared memory serves as a transmit queue for the external communication system [Brian, reassemble buffers 34, Fig 3; shared queue, 0142,0147].

7. As per claim 5, Brian-Payson disclose the external communication system is an Ethernet network [Brian, RDMA ethernet, 0009].

8. As per claim 6, Brian-Payson disclose said received packets comprise portions of encapsulated Ethernet packets [Brian, RDMA ethernet, 0009].

9. As per claim 7, Brian-Payson disclose said identifying to a transmission module comprises transferring to the transmission module a set of pointers identifying said reassembled communication rather than transferring to the transmission module said reassembled communication [Brian, header and tail, 0016].

10. As per claim 8, Brian-Payson disclose said identifying to a transmission module comprises identifying to the transmission module a head of said portion of said first linked list; and a tail of said portion of said first linked list [Payson link list, col 7 line s48-60].

11. As per claim 12, Brian-Payson disclose managing said linked lists for said queue pairs with a shared control; wherein each said location in said shared memory corresponds to an entry in said shared control; and wherein each entry in said shared control is configured to identify: a subsequent entry within the same linked list; and a

location in said shared memory corresponding to said subsequent entry [Payson, link list, col 7 line s48-60].

12. As per claim 13, Brian discloses A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of transferring a communication from an InfiniBand network to a communication system external

to the InfiniBand network, the method comprising:

- for each of multiple queue pairs terminating at the communication device, receiving packets [Brian, multiple queue pairs, Fig 3; termination, 0024];

- for each of said queue pairs:

- reassembling in a shared memory contents of said packets into communications to be transmitted to the external communication system [Brian, reassemble buffers 34, Fig 3; shared queue, 0142,0147]; and

- However Brian does not explicitly detail

- maintaining an associated linked list identifying locations in said shared memory in which said communications are reassembled; and

- when a communication is reassembled for a first queue pair, identifying to a transmission module a portion of a first linked list associated with said first queue pair, wherein said first linked list portion identifies shared memory locations in which said communication was reassembled.

In the same endeavor, Payson discloses an Infiniband architecture system area network including a packet context manager manages reassembly of cells into a packet context using storage memory and free list/link list memory for transmission, col 7 lines 48-60; wherein the memory is shared between ingress and egress data path, col 12 lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with shared memory and identifying the queue pair to the link list for transmission as taught by Payson into the Brian's apparatus in order to utilize the RDMA interface controller.

Doing so would provide an assist that detects when there may be an actual error in the network switch so that the fault can be corrected to thereby reduce the possibility of multiple faults occurring at the same time.

13. As per claim 14, Brian-Payson disclose as said packets are received from the InfiniBand network, queuing said contents directly into said shared memory [Brian, placed directly by OutLogic, 0065], wherein said shared memory serves as receive queues for each of said multiple queue pairs; and transmitting said reassembled communication from said shared memory, wherein said shared memory also serves as a transmit queue for the external communication system [Brian, reassemble buffers 34, Fig 3; shared queue, 0142,0147].

14. As per claim 15, Brian-Payson disclose said identifying to a transmission module comprises transferring to the transmission module a set of pointers identifying said

reassembled communication rather than transferring to the transmission module said reassembled communication [Brian, pointer, 0109; reassemble buffers 34, Fig 3; shared queue, 0142,0147].

15. As per claim 16, Brian discloses A method of storing a communication, received from an InfiniBand network, for transmission external to the InfiniBand network, the method comprising:

receiving a set of InfiniBand packets from an InfiniBand network, each said packet comprising a portion of a communication to be transmitted external to the InfiniBand network [Brian, RDMA, 0060];

storing said communication portions in a memory shared among multiple queue pairs of the InfiniBand network, including a first queue pair through which said set of InfiniBand packets is received [Brian, reassemble buffers 34, Fig 3; shared queue, 0142,0147];

However Brian does not explicitly detail

maintaining a first linked list for said first queue pair to identify locations in said memory in which said communication portions are stored; and when all of said communication portions are stored in said memory, scheduling said communication for transmission from said memory.

In the same endeavor, Payson discloses an Infiniband architecture system area network including a packet context manager manages reassembly of cells into a packet context using storage memory and free list/link list memory for transmission, col 7lines 48-60; wherein the memory is shared between ingress and egress data path, col 12 lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with shared memory and identifying the queue pair to the link list for transmission as taught by Payson into the Brian's apparatus in order to utilize the RDMA interface controller

Doing so would provide an assist that detects when there may be an actual error in the network switch so that the fault can be corrected to thereby reduce the possibility of multiple faults occurring at the same time.

16. As per claim 17, Brian-Payson disclose said storing comprises reassembling said communication portions into said communication [Brian, reassemble, 0061].

17. As per claim 18, Brian-Payson disclose said scheduling comprises: identifying to a transmission module a first entry in said first linked list corresponding to a location in said memory in which a first portion of said communication is stored; and identifying to a transmission module a final entry in said first linked list corresponding to a location in said memory in which a final portion of said communication is stored [Brian, scheduling, 0009].

18. As per claim 19, Brian-Payson disclose said InfiniBand packets comprise Send commands conveying said portions of said first communication [Brian, RDMA operation, 0021].

19. As per claim 20, Brian-Payson disclose said communication is an Ethernet packet [Brian, ethernet, 0009].

20. As per claim 23, Brian-Payson disclose said maintaining a first linked list comprises:

in a control structure, maintaining a first linked list of control entries, wherein each of said control entries except a final control entry identifies a subsequent control entry; and corresponding to said subsequent control entry, a location in said memory in which data received through said first queue pair are stored [Payson, packet context manager, link list, col 7 lines 48-60].

21. As per claim 24, Brian-Payson disclose said control structure and said memory are separate physical memory structures [Payson, packet context manager, col 7 lines 48-60].

22. As per claim 25, Brian-Payson disclose said control structure and said memory employ common addressing via said first linked list [Payson, packet context manager, link list, col 7 lines 48-60].

23. As per claim 26, Brian-Payson disclose said maintaining further comprises: maintaining a head pointer identifying a first control entry in said first linked list and a first location in said memory; and maintaining a tail pointer identifying said final control entry in said first linked list and a final location in said memory [Brian, header and tail, 0016].

24. As per claim 27, Brian-Payson disclose identifying to a transmission module a sub-list of said first linked list of control entries, wherein said sub-list includes control entries corresponding to all locations in said memory in which portions of said communication are stored [Payson, link list, col 7 lines 48-60].

25. As per claim 28, Brian-Payson disclose removing said sub-list of control entries from said first linked list only after said communication is transmitted [Boyd, update table, 0144].

26. As per claim 29, Brian discloses A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of storing a communication, received from an InfiniBand network, for transmission external to the InfiniBand network, the method comprising:
receiving a set of InfiniBand packets from an InfiniBand network, each said packet comprising a portion of a communication to be transmitted external to the InfiniBand network [Brian, RDMA, 0002];
storing said communication portions in a memory shared among multiple queue pairs of the InfiniBand network, including a first queue pair through which said set of InfiniBand packets is received [Brian, queue pair, 0019];

However Brian does not explicitly detail
maintaining a first linked list for said first queue pair to identify locations in said memory in which said communication portions are stored; and when all of said communication portions are stored in said memory, scheduling said communication for transmission from said memory.

In the same endeavor, Payson discloses an Infiniband architecture system area network including a packet context manager manages reassembly of cells into a packet context using storage memory and free list/link list memory for transmission, col 7lines

Art Unit: 2616

48-60; wherein the memory is shared between ingress and egress data path, col 12 lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with shared memory and identifying the queue pair to the link list for transmission as taught by Payson into the Brian's apparatus in order to utilize the RDMA interface controller.

Doing so would provide an assist that detects when there may be an actual error in the network switch so that the fault can be corrected to thereby reduce the possibility of multiple faults occurring at the same time.

27. As per claim 30, Brian-Payson disclose said scheduling comprises: identifying to a transmission module a first entry in said first linked list corresponding to a location in said memory in which a first portion of said communication is stored; and identifying to a transmission module a final entry in said first linked list corresponding to a location in said memory in which a final portion of said communication is stored [Payson, link list, col 7 lines 48-60].

28. As per claim 34, Brian discloses A computer readable medium containing a data structure configured for simultaneously queuing contents of packets as they are received from an InfiniBand network and storing communications reassembled from said contents for transmission external to the InfiniBand network, the data structure comprising:

a shared memory [Brian, a CQ can be shared, 0147], comprising:

for each of a plurality of InfiniBand communication connections, memory buckets for storing contents of packets received on the communication connection [Brian, RDMA controller, 0011];

wherein said contents of said packets are reassembled into communications in said memory buckets as said contents are stored [Brian, reassemble, 0061]

However Brian does not explicitly detail

a shared control, comprising for each of said communication connections, a linked list of control entries, wherein each said control entry is configured to identify:

a subsequent control entry in said linked list; and in said shared memory, a memory bucket corresponding to said subsequent control entry; wherein said shared memory and said shared control are configured for access by an InfiniBand receive module configured to receive said packets; and a transmit module configured to transmit said communications external to the InfiniBand network.

In the same endeavor, Payson discloses an Infiniband architecture system area network including a packet context manager manages reassembly of cells into a packet context using storage memory and free list/link list memory for transmission, col 7lines 48-60; wherein the memory is shared between ingress and egress data path, col 12 lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with shared memory and identifying the queue pair to the link list which provides a shared control for transmission as taught by Payson into the Brian's apparatus in order to utilize the RDMA interface controller.

Doing so would provide an assist that detects when there may be an actual error in the network switch so that the fault can be corrected to thereby reduce the possibility of multiple faults occurring at the same time.

29. As per claim 35, Brian-Payson disclose said shared control further comprises: a linked list of free entries, wherein each said free entry is configured to identify, in said shared memory, a free memory bucket [Payson, link list, col 7 lines 48-60].

30. As per claim 36, Brian-Payson disclose said shared control is protected by one or more of: separating bits of the structure to prevent double bit errors; and providing single error correct and double error detect protection for one or more control entries in the shared control [Payson, multiple faults occurring at the same time, col 2 lines 1-8].

31. As per claim 37, Brian-Payson disclose a control entry in said shared control is further protected by extending the single error correct and double error detect protection to include the identity of a memory bucket in said shared memory [Payson, multiple faults, col 2 line 6].

32. As per claim 38 Brian discloses An apparatus for storing data received from an InfiniBand network, for transfer to an external communication system, comprising:
an InfiniBand receive module configured to receive packets from a plurality of InfiniBand communication connections;
a transmit module configured to transmit communications to a communication system external to the InfiniBand network via one or more outbound ports [Brian, reassemble buffers 34, Fig 3];

a memory shared between said InfiniBand receive module and said transmit module
[Brian, shared queue, 0142,0147];

However Brian does not explicitly detail
a control, shared between said InfiniBand receive module and said transmit
module, for maintaining a linked list for each of said communication connections and for
each of the outbound ports.

In the same endeavor, Payson discloses an Infiniband architecture system area
network including a packet context manager manages reassembly of cells into a packet
context using storage memory and free list/link list memory for transmission, col 7lines
48-60; wherein the memory is shared between ingress and egress data path, col 12
lines 41-54]

Therefore it would have been obvious to an ordinary skill in the art at the time the
invention was made to incorporate the link list with shared memory and identifying the
queue pair to the link list for transmission as taught by Payson into the Brian's
apparatus in order to utilize the RDMA interface controller.

Doing so would provide an assist that detects when there may be an actual error
in the network switch so that the fault can be corrected to thereby reduce the possibility
of multiple faults occurring at the same time.

33. As per claim 42, wherein said communication connections are queue pairs
[Brian, queue pair, 0019].

34. As per claim 43, wherein said communication connections are virtual lanes
[Brian, virtual address, 0025].

35. As per claim 44, wherein the external communication system comprises an Ethernet network [Brian, Ethernet, 0009].

36. As per claim 45, wherein the external communication system comprises a SONET network as a design choice.

Claims 9-11,21-22,31-33,39-41 are rejected under 35 U.S.C. 103(a) as being unpatentable over Brian et al [Brian 2005/0129039 A1] in view of Payson [7,194,661 B1] and further in view of Boyd et al [Boyd 2006/0259644 A1].

37. As per claim 9, Brian-Payson disclose said maintaining an associated linked list for said first queue pair [Payson link list, col 7 line s48-60]:

However Brian-Payson does not explicitly detail

maintaining a head pointer configured to identify: a first location in said shared memory in which contents of a first packet are stored; and a first entry in a shared control structure, said first entry corresponding to said first location in said shared memory; and maintaining a tail pointer configured to identify:

a final location in said shared memory in which contents of a final packet are stored; and a final entry in said shared control structure; wherein each entry in said shared control structure that is part of said first linked list, except for said final entry, identifies a subsequent entry in said shared control structure and identifies a location in said shared memory corresponding to said subsequent entry.

In the same endeavor, Boyd discloses the InfiniBand with RDMA including a link list with head pointer and tail pointer [Boyd, 0129; 0132] and shared control structure or IP Suite Offload Engine [Boyd, 0146]

Therefore it would have been obvious to an ordinary skill in the art at the time the invention was made to incorporate the link list with head pointer, tail pointer and a shared control structure as taught by Boyd into the Brian-Payson's apparatus in order to utilize the InfiniBand-RDMA environment.

Doing so would provide a mechanism which an IPSOE can be shared between virtual hosts of a single physical host.

38. As per claim 10, Brian-Payson-Boyd disclose protecting said shared control structure by one or more of separating bits of the structure to prevent double bit errors; and providing single error correct and double error detect protection for one or more entries in the control structure [Payson, multiple faults occurring at the same time, col 2 lines 1-8].

39. As per claim 11, Brian-Payson-Boyd disclose extending the single error correct and double error detect protection to include said location in said shared memory [Payson, multiple faults occurring at the same time, col 2 lines 1-8].

40. As per claim 21, Brian-Payson-Boyd disclose said maintaining a first linked list comprises:

maintaining a head pointer identifying: a first entry in a control structure; and
a first location in said memory in which a first portion of said communication is stored;
and maintaining a tail pointer identifying: a final entry in said control structure; and

a final location in said memory in which a final portion of said communication is stored; wherein said first entry is linked to said final entry by zero or more intermediate entries in said control structure, each said intermediate entry corresponding to an intermediate location in said memory in which a portion of said communication is stored [Boyd, a link list with head pointer and tail pointer 0129; 0132].

41. As per claim 22 Brian-Payson-Boyd disclose updating said first linked list to remove said first entry, said final entry and said intermediate entries from said first linked list only after said communication is transmitted [Boyd, update table, 0144].

42. As per claim 31, Brian-Payson-Boyd disclose said maintaining a first linked list comprises:

in a control structure, maintaining a first linked list of control entries, wherein each of said control entries except a final control entry identifies: a subsequent control entry; and corresponding to said subsequent control entry, a location in said memory in which data received through said first queue pair are stored; maintaining a head pointer identifying a first control entry in said first linked list and a first location in said memory; and maintaining a tail pointer identifying said final control entry in said first linked list and a final location in said memory [Boyd, a link list with head pointer and tail pointer 0129; 0132].

43. As per claim 32, Brian-Payson-Boyd disclose identifying to a transmission module a sub-list of said first linked list of control entries, wherein said sub-list includes control entries corresponding to all locations in said memory in which portions of said communication are stored as inherent of link list.

Art Unit: 2616

44. As per claim 33, Brian-Payson-Boyd disclose removing said sub-list of control entries from said first linked list only after said communication is transmitted [Boyd, update table, 0144].
45. As per claim 39, Brian-Payson-Boyd disclose for each said communication connection a head pointer identifying a head of said linked list for said communication connection; and a tail pointer identifying a tail of said linked list for said communication connection [Boyd, a link list with head pointer and tail pointer 0129; 0132].
46. As per claim 40, Brian-Payson-Boyd disclose for each of the outbound ports a head pointer identifying a head of said linked list for the outbound port; and a tail pointer identifying a tail of said linked list for the outbound port [Boyd, a link list with head pointer and tail pointer 0129; 0132].
47. As per claim 41, Brian-Payson-Boyd disclose for each outbound queue of each of the outbound ports a head pointer identifying a head of said linked list for the outbound queue; and a tail pointer identifying a tail of said linked list for the outbound queue [Boyd, a link list with head pointer and tail pointer 0129; 0132].

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Thong H. Vu whose telephone number is 571-272-3904. The examiner can normally be reached on 6:00-3:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, *Lynn Feild* can be reached on 571-272-2092. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2616

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Thong Vu
Primary Examiner

A handwritten signature in black ink, appearing to read 'Thong Vu', with a stylized flourish extending from the end.

THONG VU
PRIMARY PATENT EXAMINER